Model-based Testing of the Java network API

Cyrille Artho
School of Computer Science and Communication
KTH, Stockholm, Sweden
Information Technology Research Institute
AIST, Osaka, Japan

Guillaume Rousset
University of Nantes
Nantes, France

Testing networked systems is challenging. The client or server side cannot be tested by itself. We present a solution using tool "Modbat" that generates test cases for Java's network library java.nio, where we test both blocking and non-blocking network functions. Our test model can dynamically simulate actions in multiple worker and client threads, thanks to a carefully orchestrated design that covers non-determinism while ensuring progress.

1 Introduction

Model-based testing derives concrete test cases from an abstract test model [5, 17]. Tools derive test sequences from a model that specifies possible test executions and expected results [1, 7, 8, 13, 17]. We use Modbat [1] because it offers an embedded domain-specific language [19] that combines extended finite-state machines [6] with low-level monitoring code written in Scala [14].

Our work on testing Java's non-blocking networking (package `java.nio`) extends previous work where we tested a custom version of that library and found several hidden defects [3]. That custom library was designed to be compatible with Java Pathfinder (JPF) [18] by working with JPF extension `net-iocache` to backtrack the effects of network input/output [10]. It stores the effects of network operations in memory, and replays them from previously stored data after backtracking.

In `java.nio`, input/output actions can be blocking or non-blocking [16]. *Blocking* actions suspend the active thread until the result is directly returned to that thread. *Non-blocking* actions return immediately, but the result may be incomplete, requiring another call later.

Our contributions are as follows:

- We simulate possible concurrency on the server side (worker threads) and multiple client sessions in parallel, and ensure a proper orchestration of all activities.
- We show how to model blocking and non-blocking functions in Java's network library `java.nio`, where non-blocking functions may return an incomplete result.

This paper is organized as follows: Section 2 gives the background of this work. Section 3 describes our client/server model for the Java network library. Section 4 concludes and outlines future work.

2 Background

2.1 Modbat

Modbat provides an embedded domain-specific language [19] based on Scala [14] to model test executions in complex systems succinctly [1]. System behavior is described using extended finite-state machines (EFSMs) [6]. An EFSM is a finite-state machine that is extended with variables, enabling

J. Kofroň, J. Tumova (Eds.): Formal Engineering approaches
to Software Components and Architectures (FESCA'17)
EPTCS 245, 2017, pp. 46–51, doi:10.4204/EPTCS.245.4

© C. Artho & G. Rousset

functions (preconditions), and update functions (actions) for each transition. Results of actions on the system under test (SUT) can be checked using assertions inside the update function.

Test cases are derived by exploring available transitions, starting from the initial state. A test case continues until a configurable limit is hit or a property is violated. Properties include unexpected exceptions and assertion failures. Assertions encode requirements, typically safety properties, and are used to check the result of a function call within a model transition. Modbat also supports exceptions and non-deterministic outcomes: If an exception or unexpected result occurs during a transition, its target state can be overridden with a different (exceptional) state [1].

Finally, Modbat offers a `launch` function, which initializes a new child model. If multiple models are active at the same time, they are executed using an interleaving semantics.

2.2 The Java Network Library

Java offers non-blocking input/output (I/O) over TCP/IP network sockets as part of the `java.nio` package [15]. Two components are essential for this work:

1. *Channels* represent connection entities. These include server-side ports that can accept an incoming connection (`ServerSocketChannel`) and connection handles to send and receive data over an active connection (`SocketChannel`).
2. *Selectors* can query multiple channels at once on their availability, chosen by using *selection keys.*

Blocking calls suspend the active thread until the complete result is returned; non-blocking calls return immediately, but with a possibly incomplete result. The application programming interface (API) of `java.nio` allows switching between blocking and non-blocking modes at any time.

2.3 Related Work

Unit testing experienced a widespread rise in software development in the late 1990s [11]. While unit testing automates test execution, model-based testing automates test design [5, 17]. Instead of designing individual test cases, test models describe entire sets of possible tests. More test tools than can be described here exist, based on state machines [1, 8, 17] or constraint specifications [7, 13]. Test models (as well as unit tests) are usually designed based on the specification [17].

For the Java API, related work [12] presents a systematic formulation of the specification of the Java API, and a run-time monitor implementation which monitors the correct usage of the API. While our work implements a monitor which verifies the correctness of implementation of the API itself (rather than its usage by an application), systematic ways to construct the specification of the API greatly interest us.

3 Test Model for the Java Network API

We start with the overall organization of our test model and add selector-based, non-blocking I/O later.

3.1 Minimalist Model for Client/Server Connections

A server using blocking I/O uses multiple threads to handle multiple connections at the same time: A server *main* thread accepts incoming requests and then spawns a *worker* thread, which uses a given connection handle and deals with the request. Each client is typically an independent process that first

```
val ssc = ServerSocketChannel.open()
"open" -> "bound" := ssc.socket().bind(...)
"bound" -> "bound" := {
  launch(new ClientModel)
  val conn = ssc.accept()
  launch(new Worker(conn)) }
"bound" -> "closed" := ssc.close()
```

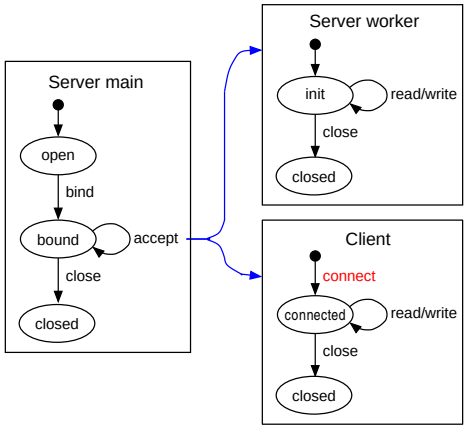

Figure 1: Minimalist model for client/server connections, with the server main model code (left) and a visualization of all three components (right). The second transition first launches a client model, then calls "accept" on the `ServerSocketChannel`, and finally passes the handle for the incoming connection to a newly launched server worker model.

connects to the server and then communicates with it. For the purpose of testing, we launch client models as child instances, which simulate external processes while being executed inside the test harness.

Figure 1 shows this parallel composition of dynamically instantiated models. It mirrors a server that accepts incoming client connections and then delegates the connection handle to a worker thread. The crucial part is transition "bound" → "bound". To ensure a correct handshake that establishes a working connection, a client model that connects back to the server in its constructor has to be launched first. The call to `connect` is executed in the constructor of the model; this ensures that the subsequent call to `accept` succeeds (because the operating system queues the pending client connection request). After that, the connection handle is passed to a new server worker model instance, which uses that connection.

3.2 Server Implementation Using Selectors

High-performance servers use non-blocking, selector-based I/O, to handle many connections in a single execution thread [9]. The server usually calls `select` in an infinite loop, which returns a set of available channels from which data can be consumed [16, 4]. Available operations include `accept`, to handle a new incoming connection, and `read`, to handle new data on an existing connection. The complexity of concurrency in the previous architecture is replaced by non-determinism w. r. t. the outcome of each `select` call and possibly incomplete `read` and `write` operations that require careful buffer management.

3.3 Server Main and Worker Models Using Selectors

Our Modbat model for `java.nio` expands on an earlier version, which was used to uncover defects in a custom version of the `java.nio` library but did not use multiple connections [3].

The detailed server main model in Figure 2 refines the original simple model with selectors and non-blocking calls, which have to be repeated until they succeed. It starts with a fully initialized instance of `ServerSocketChannel`. Each state has several self-transitions, which do not affect the state of the model or the corresponding object that is tested. Self-transitions may change between blocking and non-blocking mode (toggleBlocking), check the state of the selector (checkSelector) or the local port number

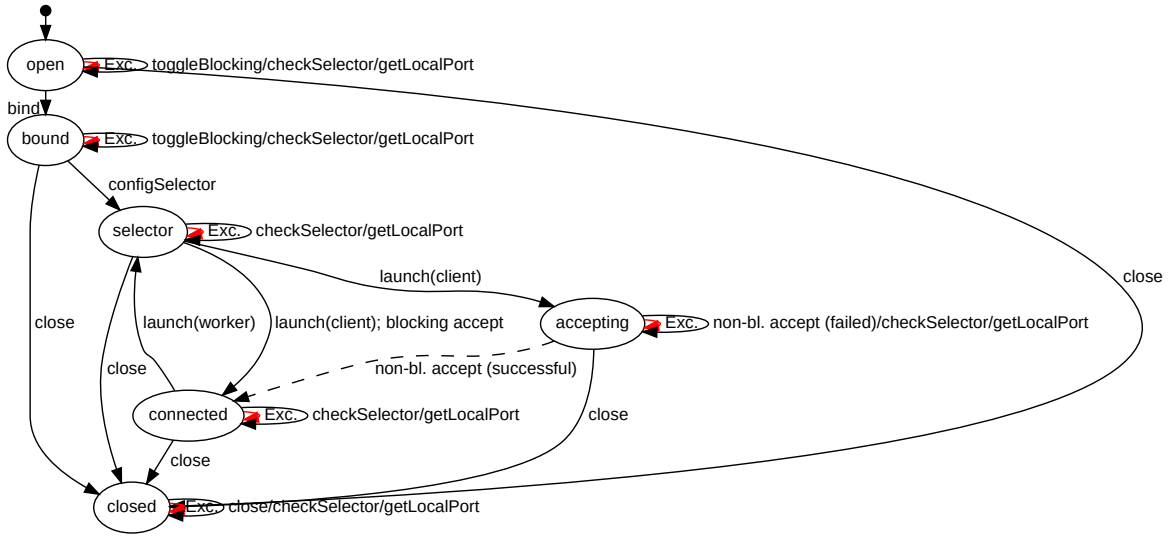

Figure 2: Detailed model of the server main thread. Labeled nodes are states, black solid arrows are model transitions. The dashed transition represents a successful non-blocking accept call, which is modeled as a non-deterministic outcome. Red arrows represent exceptions for operations that are disallowed in some states. '/" shows alternatives for a transition, ";" a concatenation of two actions.

(getLocalPort), or perform an operation that is not allowed in a given state and that results in an exception (red transitions labeled "Exc."). A "successful" path through the test model first binds the channel to an address and network port, then configures the selector, and then accepts one or more client connections before shutting down the service by calling `close` on the `ServerSocketChannel`. Calls to `accept` in non-blocking mode may initially fail because the client model may not be ready when the call is made. In that case, the model stays in an intermediate "accepting" state. In that state, the call to `accept` can be repeated until the operation is successful. As the outcome depends on network latency, a successful result is modeled as a non-deterministic outcome in Modbat, shown as a dotted transition in Figure 2. The model proceeds to state "connected" if the returned connection object is initialized (non-null) and stays in state "accepting" otherwise.

Because selector and read/write calls can be made on each connection independently, a separate "worker" model simulates such calls (see Figure 3a). In this model, the generated model instances simulate interleaved selector usages.Each connection model starts with an initialized connection, from which at can proceed by shutting down the input or output channel, or both by closing the connection. In each state, self-transitions can read or write on the channel, or check the state of its selectors. Once a channel is partially or fully closed, several operations result in an exception.

Each server-side connection model instance is paired with a minimalist client instance (see Figure 3b). This ensures a connection to the server socket is eventually established, which is necessary for testing networked software components, as they cannot be executed in isolation.

We check that the return values and exceptions thrown are consistent with the model-side view of the overall system. Due to network latency, it may be possible that data is not available in cases where an ideal network could provide data; our assertions take this into account. For example, the number of bytes read may be less or equal than the number of bytes available in the channel. The opposite case, a return value suggesting the availability of data where no data is believed (by the model) to exist, results in a property violation that is shown as an error trace by Modbat [3].

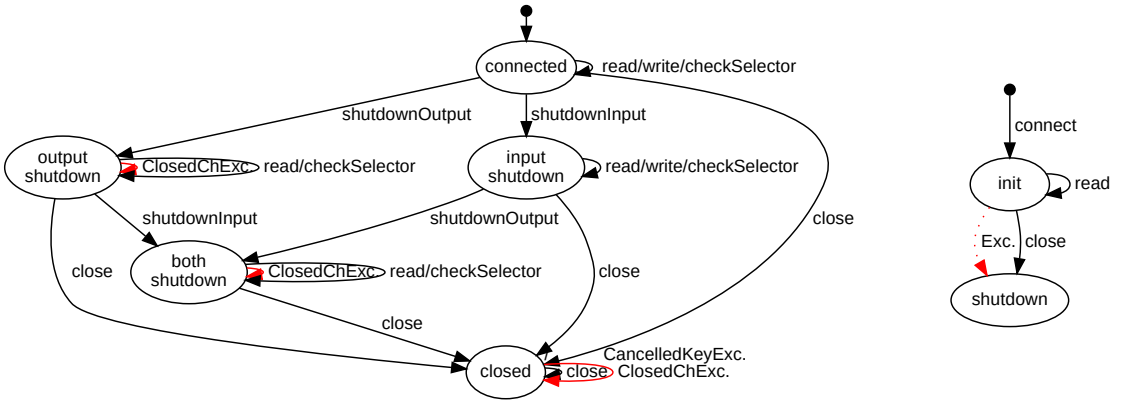

(a) Server-side connection model (b) Client model

Figure 3: Test models representing a server connection (“worker”, left) and a client connection (right). Selector calls are possible as long as the channel is not closed. The client model reads until either an exception occurs, or by a non-deterministic model choice closes the connection.

3.4 Challenges

Side effects of test actions pose a challenge. If a test does not clean up all resources when it ends, dependencies between tests may arise. We observed this in cases where a test could not be replayed in isolation but only if it was part of a larger sequence of multiple tests. We eventually arrived at what we think is a correct model by executing half a million test cases against the system, without false positives.

Test cases for the Java network API execute quickly, but ephemeral ports that handle a connection are exhausted after a while. On today’s systems, a few tens of thousands of such ports are available; their number can be increased only slightly with a custom kernel configuration. A lack of available ephemeral ports results in a slowdown of test execution (until closed ports are made available again).

In the model described in this paper, the behavior of each connection is independent of other connections, so the test oracle is entirely local to each model instance, and independent of interleavings of messages sent over the network. If this is not the case, the oracle has to consider all possible interleavings of events, as described in other work [2].

4 Conclusion and Future Work

We present a test model for a network API that uses multiple parallel model instances to simulate concurrent requests. The critical issue is to ensure progress after launching a new model instance, or when waiting for a request. To achieve this, (1) the server side must be ready to receive requests; (2) a client session is initiated; and (3), in the case of the Java network API, the necessary server call to accept the client session must be executed. Changing the order of these steps leads to a deadlock in test execution.

Our model faithfully reflects all key operations of the Java network library [16], and can be used to analyze different implementations in depth; previous work reports on defects found with our approach [3].

The models we present can be adapted to other client/server systems. In the future, we want to explore more network APIs, and other uses of Modbat as an event simulator.

Modbat and example models are available at `http://people.kth.se/~artho/modbat/`.

References

[1] C. Artho, A. Biere, M. Hagiya, E. Platon, M. Seidl, Y. Tanabe & M. Yamamoto (2013): *Modbat: A Model-based API Tester for Event-driven Systems*. In: *Proc. 9th Haifa Verification Conf.*, *LNCS* 8244, Springer, Haifa, Israel, pp. 112–128, doi:10.1007/978-3-319-03077-7_8.

[2] C. Artho, Q. Gros, G. Rousset, K. Banzai, L. Ma, T. Kitamura, M. Hagiya, Y. Tanabe & M. Yamamoto (2017): *Model-based API Testing of Apache ZooKeeper*. In: *10th IEEE International Conference on Software Testing, Verification and Validation (ICST 2017)*, IEEE, Tokyo, Japan.

[3] C. Artho, M. Hagiya, R. Potter, Y. Tanabe, F. Weitl & M. Yamamoto (2013): *Software Model Checking for Distributed Systems with Selector-Based, Non-Blocking Communication*. In: *Proc. 28th Int. Conf. on Automated Software Engineering*, ASE, IEEE, Palo Alto, USA, pp. 169–179, doi:10.1109/ASE.2013.6693077.

[4] E. Baeldung (2016): *Introduction to the Java NIO Selector*. http://www.baeldung.com/java-nio-selector. Accessed: 2017-03-07.

[5] R. Binder (2000): *Testing object-oriented systems: models, patterns, and tools*. Addison-Wesley.

[6] K. Cheng & A. Krishnakumar (1993): *Automatic functional test generation using the extended finite state machine model*. In: *Proc. 30th Int. Design Automation Conf.*, DAC, ACM, Dallas, USA, pp. 86–91, doi:10.1145/157485.164585.

[7] K. Claessen & J. Hughes (2000): *QuickCheck: a lightweight tool for random testing of Haskell programs*. *SIGPLAN Not.* 35(9), pp. 268–279, doi:10.1145/357766.351266.

[8] J. Jacky, M. Veanes, C. Campbell & W. Schulte (2007): *Model-Based Software Testing and Analysis with C#*, 1st edition. Cambridge University Press, doi:10.1017/CBO9780511619540.

[9] D. Kegel (2013): *The C10K problem*. http://www.kegel.com/c10k.html.

[10] W. Leungwattanakit, C. Artho, M. Hagiya, Y. Tanabe, M. Yamamoto & K. Takahashi (2014): *Modular Software Model Checking for Distributed Systems*. *IEEE Transactions on Software Engineering* 40(5), pp. 483–501, doi:10.1109/TSE.2013.49.

[11] J. Link & P. Fröhlich (2003): *Unit Testing in Java: How Tests Drive the Code*. Morgan Kaufmann Publishers, Inc.

[12] Q. Luo, Y. Zhang, C. Lee, D. Jin, P. Meredith, T. Şerbănuţă & G. Roşu (2014): *RV-Monitor: Efficient Parametric Runtime Verification with Simultaneous Properties*. In: *Proc. 5th Int. Conf. on Runtime Verification*, *RV* 8734, Springer, Toronto, Canada, pp. 285–300, doi:10.1007/978-3-319-11164-3_24.

[13] R. Nils (2013): *ScalaCheck, a powerful tool for automatic unit testing*. https://github.com/rickynils/scalacheck/. Accessed: 2016-12-30.

[14] M. Odersky, L. Spoon & B. Venners (2010): *Programming in Scala: A Comprehensive Step-by-step Guide*, 2nd edition. Artima Inc., USA.

[15] Oracle (2016): *Java Platform SE 8*. http://docs.oracle.com/javase/8/docs/api/.

[16] Oracle (2016): *Java Platform Standard Edition 8 API Specification*. http://docs.oracle.com/javase/8/docs/api/.

[17] M. Utting & B. Legeard (2006): *Practical Model-Based Testing: A Tools Approach*. Morgan Kaufmann Publishers, Inc., San Francisco, USA.

[18] W. Visser, K. Havelund, G. Brat, S. Park & F. Lerda (2003): *Model checking programs*. *Automated Software Engineering Journal* 10(2), pp. 203–232, doi:10.1023/A:1022920129859.

[19] D. Wampler & A. Payne (2009): *Programming Scala*. O'Reilly Series, O'Reilly Media.